DSN 2010: 1st International Workshop on Fault-Tolerance for HPC at Extreme Scale

Chicago, Illinois, USA

Call For Papers

Objectives and Challenges

With the emergence of many-core processors, accelerators, and alternative/heterogeneous architectures, the HPC community faces a new challenge: a scaling in number of processing elements that supersedes the historical trend of scaling in processor frequencies. The attendant increase in system complexity has first-order implications for fault tolerance. Mounting evidence invalidates traditional assumptions of HPC fault tolerance: faults are increasingly multiple-point instead of single-point and interdependent instead of independent; silent failures and silent data corruption are no longer rare enough to discount; stabilization time consumes a larger fraction of useful system lifetime, with failure rates projected to exceed one per hour on the largest systems; and application interrupt rates are apparently diverging from system failure rates.

The workshop will convene a diverse group of experts in HPC and fault-tolerance to inaugurate a fault-tolerance research agenda for responding to the unique challenges that extreme scale and complexity. Innovation is encouraged and discussion of non-traditional approaches is welcome.

Topics

Assuming hardware and software errors will be inescapable at extreme scale, this workshop will consider aspects of fault tolerance peculiar to extreme scale that include, but are not limited to:

- Quantitative assessments of cost in terms of power, performance, and resource impacts of fault-tolerant techniques, such as checkpoint restart, that are redundant in space, time or information;
- Novel fault-tolerance techniques and implementations of emerging hardware and software technologies that guard against silent data corruption (SDC) in memory, logic, and storage and provide end-to-end data integrity for running applications;
- Studies of hardware / software tradeoffs in error detection, failure prediction, error preemption, and recovery;
- Advances in monitoring, analysis, and control of highly complex systems;
- Highly scalable fault-tolerant programming models;
- Metrics and standards for measuring, improving and enforcing the need for and effectiveness of fault-tolerance;
- Failure modeling and scalable methods of reliability, availability, performability and failure prediction for fault-tolerant HPC systems;
- Scalable Byzantine fault tolerance and security from single-fault and fail-silent violations;
- Benchmarks and experimental environments, including fault-injection and accelerated lifetime testing, for evaluating performance of resilience techniques under stress.

Participation and Paper Submission

Submissions are expected in the following categories:

- Extended abstracts that propose original ideas in the field;
- Work-in-progress reports that present considerable progress in the challenging areas;
- Position papers that identify open issues or discuss existing solutions.

The submissions shall be sent electronically, must conform to IEEE conference proceedings style and should not exceed six pages including all text, appendices, and figures.

Important Dates

Submission of papers:	March 15, 2010
Author notification:	April 9, 2010
Camera ready papers:	April 30, 2010

Further Information

http://institute.lanl.gov/resilience/conferences/ftxs2010/

Workshop location, registration and accommodation: http://www.dsn.org